

Long-Wire Leakage: The Threat of Crosstalk

Ilias Giechaskiel, Kasper B. Rasmussen, and Ken Eguro

in *IEEE Design & Test*, vol. 39, no. 4, pp. 41–48, Aug. 2022.

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
DOI:10.1109/MDAT.2022.3142199

Long-Wire Leakage: The Threat of Crosstalk

Ilias Giechaskiel
Independent Researcher
London, UK
iliask@giechaskiel.com

Kasper Rasmussen
University of Oxford
Oxford, UK
kasper.rasmussen@cs.ox.ac.uk

Ken Eguro
Microsoft Corporation
Redmond, WA, USA
eguro@microsoft.com

Abstract—With Field-Programmable Gate Arrays (FPGAs) becoming larger, heterogeneous, and more widely available in data center environments, it is important to consider how low-level hardware choices can affect the security of user designs. An example of an electrical effect that can compromise sensitive data through covert- and side-channel attacks is capacitive crosstalk between long routing wires connecting logic resources that are physically far apart within the same Integrated Circuit (IC). This article summarizes recent developments showing that this novel source of information leakage is present on Xilinx and Intel FPGAs, as well as Application-Specific Integrated Circuits (ASICs). It can be exploited in devices spanning several technology nodes and architectures, does not require physical access to the FPGA board, and can be measured using just on-chip resources. This article further presents existing software- and hardware-based defense mechanisms, and identifies open questions and future research directions. Overall, this article highlights the shift in the system and adversary model used in analyzing hardware security, and therefore the need for new IC designs to incorporate countermeasures protecting against the threats arising from multi-tenant occupancy of chips.

Index Terms—Long-wire delays; crosstalk; information leakage; routing resources; FPGA security; FPGA remote attacks; side channels; covert channels; fault attacks

I. INTRODUCTION

THE ability of Field-Programmable Gate Arrays (FPGAs) to implement highly-parallel, reconfigurable hardware makes them popular in consumer end-products [1], cloud accelerators for cryptography, genomic sequencing, and financial modeling [2], or even in military and aerospace applications, such as radars and electronic warfare [3]. In parallel, the ever-increasing complexity of FPGA designs (coupled with the recent availability of FPGAs on the cloud, and the heterogeneous nature of high-end, multi-die chips) has led to a new threat model, one of *remote* attacks to FPGAs without physical access to the device or external equipment such as high-end oscilloscopes and measurement probes.

Indeed, ever since Amazon Web Services (AWS) announced F1 instances with FPGAs on their Elastic Cloud Compute (EC2) platform [4], multi-tenant and virtualized [5], [6] FPGAs have come closer to being a reality. These setups aim to better utilize the large amount of parallelizable hardware resources (e.g., 1.3 million lookup tables spread over three dies for Xilinx VU9P chips [7]) by sharing them among multiple users. Users have access to only parts of the reconfigurable hardware through partial reconfiguration, with a portion of the FPGA dedicated to a cloud-provided “shell”. The shared

on- and off-chip infrastructure (e.g., common Power Supply Units [8] or PCIe bandwidth contention [9]) can result in information leakage that is exploited by malicious Intellectual Property (IP) cores or adversarial users for covert- and side-channel attacks, in effect breaking the security guarantees that the logical isolation of different tenants aims to achieve.

One of the earliest works recognizing the potential for unintentional interactions between different electrical elements in FPGAs showed that a certain type of routing resource, called a *long* wire, can leak information about its state to adversarial circuits by influencing the delays of nearby long wires [10]. In the years since, follow-up work has demonstrated that long-wire leakage presents a security concern in Xilinx [7], [11]–[13] and Intel [14], [15] FPGAs spanning over a decade of different technology nodes and architectures, as well as Application-Specific Integrated Circuits (ASICs) [16]. Using just on-chip sensors called *Ring Oscillators* (ROs), it can be exploited for covert-channel attacks with high-bandwidth and accuracy (over 6 kbps with fewer than 0.1% errors) [11], or even in side-channel attacks for recovering AES keys in logic running at 10 MHz through repeated measurements [14].

This article identifies the key findings of the literature in the area, synthesizing information about trends and the root cause of the phenomenon (Section II), presents existing proposals for software- and hardware-based defense mechanisms (Section III), and concludes by discussing open questions and future research directions (Section IV).

II. LONG-WIRE LEAKAGE

Work on long-wire leakage features prominently in surveys on FPGA attacks and defenses [17]–[19], in part because it was one of the first sources of information leakage to be demonstrated in the multi-tenant system and adversary model, where logically isolated blocks are able to infer information about each other, even when they are not directly connected. However, unlike later work which depends on temperature changes [20] or voltage drops [21]–[24], long-wire leakage persists even when the driven value remains constant. Static leakage is known to be harder to detect, requiring large circuits, measurement intervals, external equipment, and special modifications to the device under test [25]. By contrast, the state of a long wire can be detected by small, on-chip circuits, without physical access to or modification of the FPGA board.

This section summarizes existing work on characterizing long-wire leakage by first describing how it is measured




Fig. 1. Measurement setup for long-wire leakage: the attacker attempts to infer the values carried by the victim long wires by measuring the frequency of a ring oscillator routed through nearby long wires. The attacker and victim circuits are logically isolated.

(Section II-A) on Xilinx (Section II-B) and Intel (Section II-C) FPGAs, and then discussing similar results on ASICs which point towards the root cause of the phenomenon (Section II-D).

A. Measurement Setup

Most experiments on measuring long-wire leakage use the same basic setup, shown in Figure 1: a ring oscillator is constructed by chaining an odd number of NOT gates in a combinatorial loop, with two of its stages connected using one or more long wires. These wires are chosen so that they are routed adjacent to one or more “victim” long wires, whose value the adversary aims to infer. The attacker and victim circuits are otherwise entirely independent, i.e., logically isolated. The ring oscillator then drives a counter, which is sampled at a fixed measurement interval. When the victim long wires carry a logic 1, the delays of attacker long wires become slightly shorter compared to when the victim long wires carry a logic 0. This has the effect of increasing the frequency of the ring oscillator, and, by extension, the counter values within the measurement period also become higher.

Three metrics have been used to estimate the strength of the information leakage: (a) the absolute count difference, $\Delta C = C_{RO}^1 - C_{RO}^0$, between ring oscillator counts when the victim wire carries a logic 1 (C_{RO}^1) and a logic 0 (C_{RO}^0), (b) the relative count (or frequency) difference, $\Delta RC = \frac{C_{RO}^1 - C_{RO}^0}{C_{RO}^1} \approx \frac{f_{RO}^1 - f_{RO}^0}{f_{RO}^1}$, which is independent of the measurement duration and system clock period, and (c) the absolute delay difference, $\Delta d = \frac{C_{CLK}}{2f_{CLK}} \cdot \frac{C_{RO}^1 - C_{RO}^0}{C_{RO}^1 C_{RO}^0} \approx \frac{1}{2} \cdot \left(\frac{1}{f_{RO}^0} - \frac{1}{f_{RO}^1} \right)$, which estimates the change in the propagation delay for a signal going around the receiver ring oscillator loop by accounting for the system clock frequency f_{CLK} and the number of clock transitions during the measurement interval C_{CLK} .

B. Xilinx FPGAs

Long-wire leakage has been demonstrated on Xilinx FPGA boards from the 90 nm Virtex 4 family released in 2004 [10] to the 16 nm Virtex UltraScale+ family released in 2016 [7]. Although Giechaskiel et al. first showed in 2016 that long-wire leakage can be a security concern [10], Gag et al. had earlier established that the delay of long wires can pose a reliability issue [26]. In their 2012 paper, Gag et al. determined that the delay of a “victim” long wire is lowest when its two

TABLE I
VERTICAL LONG-WIRE PROPERTIES OF THE XILINX VIRTEX 5, VIRTEX 6, 7 SERIES, AND VIRTEX ULTRASCALE+ FAMILIES.

Property	Virtex 5	Virtex 6	7 Series	Virtex US+
Node Size	65 nm	40 nm	28 nm	16 nm
Bidirectional	Yes	Yes	Yes	No
Span (# of CLBs)	18	16	18	12
# of Read-Only Taps	2	1	1	0
# of VLONGs/CLB	2	2	2	2×8

adjacent “aggressor” long wires carry a signal that is in sync with that of the victim signal, and highest when the aggressor signals oppose it [26]. In particular, Gag et al. showed that a ring oscillator using four long wires becomes up to 9.6% faster (the frequency increases from 131.7 MHz to 144.4 MHz) in Virtex 6 devices and 7% faster in Virtex 5 devices (from 91.9 MHz to 98.4 MHz), corresponding to a delay change of 672 ps and 719 ps respectively.

The patterns tested by Gag et al. required that the signals on the aggressor and victim long wires be in sync, and were therefore not immediately applicable to covert- and side-channel attacks in the multi-tenant context: the aggressor and victim wires were directly connected to each other, and as they were driven by a common source, patterns that are independent of the ring oscillator signal were not tested. By contrast, Giechaskiel et al. [7], [10]–[12] showed that the effects of long-wire leakage are measurable for static and dynamic signals that are not under the adversary’s control. More concretely, it was shown that the ring oscillator frequency is linearly dependent on the Hamming Weight of the victim signal during the measurement period, and not the signal’s switching activity. This fact distinguished the underlying mechanism of information leakage from other effects reported in the literature (e.g., voltage drops), and was shown to be suitable for on-chip covert-channel transmissions and side-channel exfiltration attacks. A “windowing” approach with repeated measurements was shown to be capable of eavesdropping on 770 kbps signals with an accuracy of more than 98.4% [12], while a single-measurement Manchester encoding scheme resulted in a covert channel with a bandwidth of 6 kbps and 99.9% accuracy [11], even in the presence of environmental noise (local temperature and voltage fluctuations).

Giechaskiel et al. further showed that long-wire leakage is present across all absolute and relative placements of the attacker and victim long wires (provided that they remain adjacent) [11], [12], and can be measured through different types of ring oscillator receivers [7], [13]. In other words, the absolute change in the delay of a long wire does not depend on (a) the measurement period of the receiver ring oscillator, (b) which specific segment of the ring oscillator uses the long wire, (c) where the ring oscillator is placed on the device, or (d) whether the signals in the (bi-directional) long wires are propagating upwards or downwards. In fact, Giechaskiel et al. showed that by chaining multiple long wires together, the information leakage becomes stronger,




Fig. 2. Absolute delay difference (Δd) for different Xilinx boards and numbers of transmitter and receiver long wires $v_t = v_r$, synthesized using data from [7], [11]–[13] and original measurements.

TABLE II
SIZE-RELATED PROPERTIES OF THE XILINX BOARDS.

Board	Family	Part Number	LUTs	Node	Taps
ML509	Virtex 5	XC5VLX110T	69,120	65 nm	2
Basys3	Artix 7	XC7A35T	20,800	28 nm	1
ArtyS7	Spartan 7	XC7S50	32,600	28 nm	1
PYNQ	Zynq-7000	XC7Z020	53,200	28 nm	1
Nexys4	Artix 7	XC7A100T	63,400	28 nm	1
ML605	Virtex 6	XC6VLX240T	150,720	45 nm	1
KC705	Kintex 7	XC7K325T	203,800	28 nm	1
VCU118	Virtex US+	XC7VU9P	1,182,240	16 nm	0
AWS	Virtex US+	XC7VU9P	1,182,240	16 nm	0
Huawei	Virtex US+	XC7VU9P	1,182,240	16 nm	0

making it easier to discern the nearby state: as the number of overlapping transmitter and receiver long wires increases, the change in delay also increases in a proportional manner [11], [12]. Similarly, the ring oscillator can simultaneously infer information about both of its adjacent victim wires (to the left and right of the adversarial receiver): the delay is highest when both victim wires are carrying logic 0s, lowest when both are carrying logic 1s, and in between otherwise [12]. Covert-channel attacks can exploit this fact to increase the accuracy and/or the channel bandwidth through multi-bit transmissions.

Table I summarizes the properties of the Xilinx families tested, which include the 65 nm Virtex 5, the 40 nm Virtex 6, the 28 nm Spartan 7, Artix 7, Kintex 7, and Zynq-7000, as well as the 16 nm Virtex UltraScale+.¹ Prior to the Virtex UltraScale+, long wires were bi-directional, and Configurable Logic Blocks (CLBs) could drive two long wires, one towards the top of the device, and the other towards the bottom, with adjacent long wires originating from adjacent CLBs. Long

¹Most research on long-wire leakage has focused on *vertical* long wires (VLONGs), as the properties of Xilinx FPGA chips (which have tall, heterogeneous columns) make it easier to chain multiple vertical long wires than horizontal ones. However, Giechaskiel et al. noted that the phenomenon is also present in horizontal long wires [10]–[12]. For similar reasons, even though the phenomenon was shown to exist as far back as the 90 nm Virtex 4 [10], the device was deemed too small for more extensive experimentation.

wires spanned 18 CLBs (except in the Virtex 6 family, where they spanned 16 CLBs), and had intermediate read-only taps. Figure 2 presents the absolute delay difference (Δd) of the receiver ring oscillator for different (equal) numbers of chained transmitter and receiver long wires $v_t = v_r$ on several Xilinx families and boards, with results from identical devices shown in the same color. There are three main conclusions to draw from this figure: first, long-wire leakage is a fundamental problem in several FPGA generations; second, due to manufacturing variations, the strength of the information leakage is not the same between otherwise identical boards (but generally remains similar); and third, femto- to pico-second changes in the long wire delays are measurable using small on-chip circuits, and pose a security threat in multi-tenant setups.

Inspecting the size-related properties of the boards tested in Table II potentially suggests two further insights: first, the number of intermediate read-only taps appears to be significant; and second, for a given number of taps, long-wire leakage in larger devices seems to be less pronounced. Indeed, the Virtex 5 with its two intermediate taps has the strongest per-CLB leakage compared to the Virtex 6 or 7 Series with one tap, or the Virtex UltraScale+ with no taps. Though part of this effect may be due to different architectural choices and process nodes, a different type of “medium” wire in 7 Series devices, which spans only 12 CLBs and does not have intermediate taps further supports this hypothesis: according to Giechaskiel and Szefer, the per-CLB change in long-wire delays is higher than the per-CLB change in medium-wire delays, suggesting physical or electrical differences between the two types of wires [13]. Figure 3 presents the absolute delay difference (Δd) in medium wires for 7 Series devices, and includes the long-wire results for Virtex UltraScale+ FPGAs tested both locally as well as in the Amazon Web Services (AWS) and Huawei clouds. This is because Virtex UltraScale+ long wires more closely resemble 7 Series medium wires in terms of their length (they both span 12 CLBs), and their lack of intermediate read-only taps. Virtex UltraScale+ long wires are however uni-directional, and come in routing channels of eight in each direction, with adjacent wires originating from the same CLB.

C. Intel FPGAs

In 2018, Ramesh et al. showed that most of the insights for long-wire leakage on Xilinx FPGAs also hold for Intel FPGAs: C4 wires in Cyclone IV E, Cyclone IV GX, and Stratix V GX FPGAs leak information about their state in a way that is independent of the device location and measurement period [14]. Long-wire leakage is further additive in the number of chained receiver and transmitter long wires, whether the transmitter long wires are routed to the left or to the right of the ring oscillator receiver [14]. More importantly, Ramesh et al. demonstrated that long-wire leakage can be exploited in practice to extract AES key bytes through a technique similar to Differential Power Analysis (DPA). Through repeated measurements, keys can be recovered even when the AES core is running at 10 MHz, and even for auto-placed and auto-routed signals, though the number of Measurements-to-

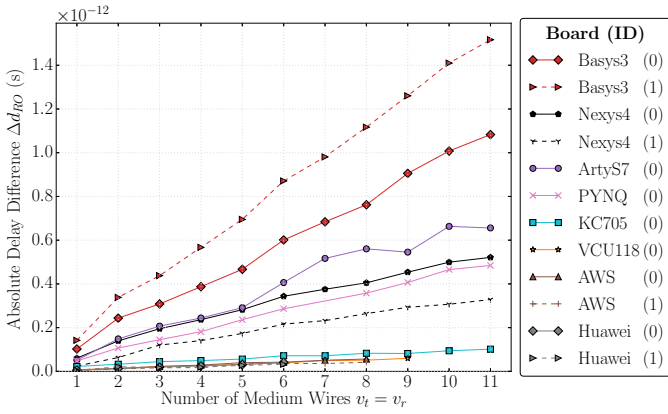


Fig. 3. Absolute delay difference (Δd) for different Xilinx boards and numbers of transmitter and receiver medium wires $v_t = v_r$, synthesized using data from [7], [11]–[13] and original measurements.

TABLE III

PROPERTIES AND LONG-WIRE DELAY CHANGES FOR INTEL BOARDS [15].

Family	LEs	Node	Δd_{C4}^L	CN	Δd_{CN}^L
Cyclone IV GX	149,760	60 nm	47.8 fs	C16	14.6 fs
Stratix V GX	622,000	28 nm	14.0 fs	C14	3.9 fs
Arria 10 GX	1,150,000	20 nm	8.2 fs	C27	16.5 fs

Disclosure (MTD) increases with the higher core frequencies and smaller overlaps between the side-channel receiver and the victim long wires [14].

A year later, Provelengios et al. reverse-engineered the routing channel layouts and more precisely characterized the changes in the delays of long wires due to adjacent state [15]. Specifically, Provelengios et al. investigated C4 wires, which span 4 Logic Array Blocks (LABs) in channels of 96 wires, half of which travel upwards, and the rest of which propagate downwards [15]. Each LAB can connect to 12 of these wires (in each direction), and Provelengios et al. measured the effect of each of the 48 C4 wires on the remaining wires in the channel to exhaustively characterize the information leakage between all possible wire pairs. For most receiver wires, there were exactly two other wires (likely the left and right neighbors) that affect the receiver wire delays, with the effect being symmetric (i.e., if wire x affects wire y , then wire y affects wire x) [15].

Armed with this information, Provelengios et al. calculated the absolute delay difference per LAB Δd_{C4}^L for the 60 nm Cyclone IV GX, the 28 nm Stratix V GX, and the 20 nm Arria 10 GX to be 47.8 fs, 14.0 fs, and 8.2 fs respectively [15]. They further calculated the absolute delay Δd_{CN}^L for a different type of wire, whose length differs per generation: Cyclone IV GX C16 wires have $\Delta d_{C16}^L = 14.6$ fs, Stratix V GX C14 wires have $\Delta d_{C14}^L = 3.9$ fs, while Arria 10 GX C27 wires have $\Delta d_{C27}^L = 16.5$ fs [15]. These results, summarized in Table III, identify that long-wire leakage is present with several wire types, and exists on Intel devices released as far back as 2009 (Cyclone IV GX) and as recently as 2013 (Arria 10 GX).

However, there is no clear trend for how long-wire leakage varies across technology nodes (even when normalizing for the process size [15]), so, as with Xilinx FPGAs, further research is needed for a better understanding of these discrepancies.

D. ASICs & Leakage Cause

Although the works on FPGA long wires have focused on the threat of information leaks from the victim wire to adversarial ones, the converse problem of injecting faults from adversarial wires into victim ones has been extensively explored in the realm of Application-Specific Integrated Circuits (ASICs) [16]. The threat model in this case applies to general-purpose processors as well as setups that integrate potentially untrusted IP or manufacturing processes, and is one of Hardware Trojans that can be implemented simply by re-routing existing resources, and without violating Design Rule Checks (DRCs). The aggressor wires push the victim voltage level past the logic threshold, therefore injecting faults to extract AES keys or escalate CPU privileges [16].

Another aspect of the 2019 work by Kison et al. that differs from that of the prior investigations on long-wire leakage is that it delved deeper into the root cause of the phenomenon. FPGA research had broadly classified the undesired parasitic effects as “crosstalk” [11], [14], [26], but could not further pinpoint the cause due to the proprietary nature of the physical layout and process parameters for the IC. By contrast, Kison et al. ran simulations using the 45 nm NanGate FreePDK45 library to model the coupling capacitance between adjacent wires [16]. Capacitive crosstalk “depends on a change in the voltage [...] that generates noise pulses”, with the coupling capacitance being proportional to the area between wires and to the inverse of the distance between them [16].

Despite some differences in the FPGA and ASIC setups, several results are consistent with or point towards capacitive crosstalk being the culprit behind long-wire leakage:

- 1) The victim signal can only be detected from an adversary at most two long wires away [11], similar to how a fifth aggressor (i.e., a third wire on a given side) “decreases reliability” due to its small contribution [16].
- 2) The change in delay increases proportionally with the number of chained long wires [11], [14], much like the coupling capacitance increases linearly with the area, which depends on the wire length [16].
- 3) Simultaneous signal transitions have similar effects in both the FPGA [26] and the ASIC [16] environments.
- 4) Concurrent transmissions from both the left and the right of the measurement ring oscillator behave identically to chained consecutive wires [12], the same way as the “multiple aggressors cause an additive effect” [16].
- 5) Intermediate read-only taps (and their purported drivers) affect the strength of the information leakage [13], and are also identified as potential areas to investigate for mitigations [16].
- 6) The magnitude of the information leakage remains the same whether the receiver and the transmitter long wires are driven in the same direction or not [11].

III. DEFENSE MECHANISMS

Besides high-level calls for architectural changes in the physical layout of long wires for future ICs (e.g., wire spacing and metal layer use [13], [16], which depend closely on technology node and foundry requirements), countermeasures for existing chips primarily require routing sensitive signals away from potentially malicious ones [11]–[13], [27], [28]. Indeed, dynamic activity alone (e.g., large adders [11] or a cloud-provided shell [7]) is not enough to hide the changes in the long-wire delays, due to the localized nature of the leakage.

For ASICs, Kison et al. developed a tool which identifies long wires and calculates the interconnect capacitance matrix for the 100 longest wires on the chip [16]. Wires for which the capacitance is higher than twice the average are marked as “suspicious”—an approach which correctly identified the inserted Hardware Trojans [16]. Kison et al. note that this approach might not detect more sophisticated attacks which violate Design Rule Checks (DRCs), alter the chemical composition of the chip, or use more complex routing paths [16].

Defending against multi-tenant attacks is a more constrained problem: users can label sensitive wires, and either avoid routing them through long-wires [27], or ensure that they are not routed next to potentially malicious ones [28]. When designs cannot avoid long wires, e.g., due to congestion, or because they connect to external I/O, frameworks can reserve the wires which are adjacent to the security-critical victim wires and set them to constant or random values [27]. Finally, an alternative approach is to prevent adversarial receiver circuits from being able to detect the changes in delays by introducing new DRCs that check bitstreams for ring oscillators or Time-to-Digital Converters (TDCs) using latches [29], [30].

IV. CONCLUSION & FUTURE OUTLOOK

With the rise of FPGAs in cloud computing, and with large, multi-die chips paving the way for multi-tenant FPGAs, there is now a larger attack surface that forces designers to consider potential information leakage from their logic to malicious on-chip adversaries. Long-wire leakage exposes a fundamental hardware issue that is present in over a decade’s worth of devices spanning several process nodes from both major FPGA manufacturers, and can be an issue even in ASICs. The femto- to pico-second changes in the delays of long wires due to adjacent state are measurable without physical access to the device or external equipment, and thus introduced the first remote covert- and side-channel attacks on FPGAs, welcoming further research in the area. Although recent works have primarily examined the effects of temperature changes and voltage drops on multi-tenant FPGAs, long-wire leakage highlighted the importance of investigating the effects of low-level hardware imperfections that can leak information even for static or slowly-changing state, which is known to be harder to detect or measure. Progress is already being made in this area, with a recent work identifying that multiplexers in Configurable Logic Blocks (CLBs) also inadvertently leak information about the signals that pass through them [13].

However, several aspects of long-wire leakage remain open areas for future research. The main question that remains unaddressed is how changes to the technology node affect the security of the FPGA ICs through long-wire leakage. In other words, a deeper understanding of the differences between vertical, horizontal, medium, and long wires within a single FPGA family, as well as the evolution of long wires throughout different physical process nodes is needed to design new chips that are immune to crosstalk effects. In parallel, further experimentation, modeling, and simulations are needed to identify whether long wires in FPGAs are susceptible to fault injection attacks, and, if so, how to extend existing countermeasures to account for them. Overall, the threat of long-wire leakage is one of the earliest examples of the threats posed by multi-tenant FPGAs, so before such setups can become a practical reality, software routing algorithms must be modified to account for it and hardware designers must address it in future chip architectures.

REFERENCES

- [1] iFixit, “iPhone 7 teardown,” <https://www.ifixit.com/Teardown/iPhone+7+Teardown/67382>, 2016.
- [2] Amazon Web Services, “The agility of F1: Accelerate your applications with custom compute power,” https://d1.awsstatic.com/Amazon_EC2_F1_Infographic.pdf, 2021.
- [3] Intel, “FPGA for military applications – Intel FPGA,” <https://www.intel.co.uk/content/www/uk/en/government/products/programmable/applications.html>, 2021.
- [4] Amazon Web Services, “Developer preview—EC2 instances (F1) with programmable hardware,” <https://aws.amazon.com/blogs/aws/developer-preview-ec2-instances-f1-with-programmable-hardware/>, 2016.
- [5] A. Khawaja, J. Landgraf, R. Prakash, M. Wei, E. Schkufza, and C. J. Rossbach, “Sharing, protection, and compatibility for reconfigurable fabric with AMORPHOS,” in *OSDI*, 2018.
- [6] A. Vaishnav, K. D. Pham, and D. Koch, “A survey on FPGA virtualization,” in *FPL*, 2018.
- [7] I. Giechaskiel, K. B. Rasmussen, and J. Szefer, “Measuring long wire leakage with ring oscillators in cloud FPGAs,” in *FPL*, 2019.
- [8] —, “C³APSULE: Cross-FPGA covert-channel attacks through power supply unit leakage,” in *S&P*, 2020.
- [9] S. Tian, I. Giechaskiel, W. Xiong, and J. Szefer, “Cloud FPGA cartography using PCIe contention,” in *FCCM*, 2021.
- [10] I. Giechaskiel and K. Eguro, “Information leakage between FPGA long wires,” <https://arxiv.org/abs/1611.08882v1>, 2016.
- [11] I. Giechaskiel, K. B. Rasmussen, and K. Eguro, “Leaky wires: Information leakage and covert communication between FPGA long wires,” in *ASIACCS*, 2018.
- [12] I. Giechaskiel, K. Eguro, and K. B. Rasmussen, “Leakier wires: Exploiting FPGA long wires for covert- and side-channel attacks,” *TRETS*, vol. 12, no. 3, Sep. 2019.
- [13] I. Giechaskiel and J. Szefer, “Information leakage from FPGA routing and logic elements,” in *ICCAD*, 2020.
- [14] C. Ramesh, S. B. Patil, S. N. Dhanuskodi, G. Provelengios, S. Pillement, D. Holcomb, and R. Tessier, “FPGA side channel attacks without physical access,” in *FCCM*, 2018.
- [15] G. Provelengios, C. Ramesh, S. B. Patil, K. Eguro, R. Tessier, and D. Holcomb, “Characterization of long wire data leakage in deep submicron FPGAs,” in *FPGA*, 2019.
- [16] K. Kison, O. M. Awad, M. Fyrbiak, and C. Paar, “Security implications of intentional capacitive crosstalk,” *TIFS*, vol. 14, no. 12, Dec. 2019.
- [17] C. Jin, V. Gohil, R. Karri, and J. Rajendran, “Security of cloud FPGAs: A survey,” <https://arxiv.org/abs/2005.04867>, 2020.
- [18] S. S. Mirzargar and M. Stojilović, “Physical side-channel attacks and covert communication on FPGAs: A survey,” in *FPL*, 2019.
- [19] J. Zhang and G. Qu, “Recent attacks and defenses on FPGA-based systems,” *TRETS*, vol. 12, no. 3, Sep. 2019.
- [20] S. Tian and J. Szefer, “Temporal thermal covert channels in cloud FPGAs,” in *FPGA*, 2019.

- [21] I. Giechaskiel, K. B. Rasmussen, and J. Szefer, "Reading between the dies: Cross-SLR covert channels on multi-tenant cloud FPGAs," in *ICCD*, 2019.
- [22] O. Glamocanin, L. Coulon, F. Regazzoni, and M. Stojilović, "Are cloud FPGAs really vulnerable to power analysis attacks?" in *DATE*, 2020.
- [23] D. Mahmoud and M. Stojilović, "Timing violation induced faults in multi-tenant FPGAs," in *DATE*, 2019.
- [24] M. Zhao and G. E. Suh, "FPGA-based remote power side-channel attacks," in *S&P*, 2018.
- [25] A. Moradi, "Side-channel leakage through static power," in *CHES*, 2014.
- [26] M. Gag, T. Wegner, A. Waschki, and D. Timmermann, "Temperature and on-chip crosstalk measurement using ring oscillators in FPGA," in *DDECS*, 2012.
- [27] Y. Luo and X. Xu, "HILL: A hardware isolation framework against information leakage on multi-tenant FPGA long-wires," in *FPT*, 2019.
- [28] Z. Seifoori, S. S. Mirzargar, and M. Stojilović, "Closing leaks: Routing against crosstalk side-channel attacks," in *FPGA*, 2020.
- [29] J. Krautter, D. R. E. Gnad, and M. B. Tahoori, "Mitigating electrical-level attacks towards secure multi-tenant FPGAs in the cloud," *TRETS*, vol. 12, no. 3, Sep. 2019.
- [30] T. M. La, K. Matas, N. Grunchevski, K. D. Pham, and D. Koch, "FP-GADefender: Malicious self-oscillator scanning for Xilinx UltraScale+ FPGAs," *TRETS*, vol. 13, no. 3, Sep. 2020.

Ilias Giechaskiel is a Hardware Research Engineer at Jump Trading. He conducts research on FPGA security as an independent researcher and as a Research Affiliate at the Computer Architecture and Security Lab of Yale University. He holds a DPhil in Cyber Security from the University of Oxford, and is an IEEE and ACM member.

Kasper Rasmussen got his Ph.D. at the Department of Computer Science at ETH Zurich. After completing his Ph.D., he worked as a post-doc at University of California Irvine, before joining the University of Oxford. He was awarded a University Research Fellowship from the Royal Society in London, and he is a Senior member of the IEEE. Kasper Rasmussen is currently Associate Professor of Computer Science at the University of Oxford where he leads a research group that works on different aspects of system and communication security, e.g., Security of Wireless Networks, Protocol design, Applied Cryptography, Security of embedded systems and Cyber-physical systems.

Ken Eguro is a Principal Researcher at Microsoft in Redmond, WA, USA and an Associate Affiliate Professor in the Electrical and Computer Engineering Department at the University of Washington in Seattle, WA, USA.